For: gray
Printed on: Wed, Nov 23, 1994 14:37:39

Document: pl
Last saved on: Wed, Nov 23, 1994 14:40:36

Incorporating Emotional Agents Into The Cooperation Mechanism Of The Accord System

Michael A. Gray and D. Patrick Studdard
Department of Computer Science and Information Systems, American University
Washington, DC 20016–8116

The Accord project is a long-term research program aimed at understanding computationally-based multiagent systems and group intelligence. The objective is a systematic, coherent, and consistent account of the nature and properties of multiagent systems and group intelligence. The approach of the project at the individual agent level is to start with theories of the properties, capabilities, and architectures of individual intelligence, and construct agents capable of participating in intelligent groups. To succeed at this task, the project must have a sufficiently powerful agent model to support the ultimate goals. The project has chosen to base that model on one described by Genesereth and Nilsson (1990).

The Accord project is investigating modifications of Genesereth and Nilsson's hysteretic agents to account for agents whose internal state model allows a representation of itself, its environment, and other agents. These agents are capable of handling changes to their models (fixed model agents are simpler to deal with, but changeable model agents are the goal of most AI research), which allow changes in their environments produced by natural cause and by other agents. Agent reasoning must be able to accept changes in current situations as well as to add new information without losing consistency. Hence, a mechanism for keeping the agent's beliefs in a consistent state must be provided in their reasoning components. A system of dependent beliefs can be represented using a justification–based reason maintenance system. Justification–based reason maintenance (Doyle, 1979) has been successfully employed in a number of reasoning systems, most notably the Proteus system (Petrie, 1986). An algorithm for achieving complete and consistent labelling of the belief base has been given by Russinoff (1985). Bridgeland and Huhns (1990) have applied justification–based reasoning technology to multiagent systems, extending the concepts of consistency and completeness to allow for beliefs distributed among a set of agents. The Accord project has adopted multiagent justification–based reason maintenance as the belief storage and revision mechanism for its agent architectures.

Cooperation in The Accord System

The Accord project is studying a notion of multiagent cooperation that consists of the use of shared belief consistency between rational agents as the means of cooperation. When rational agents share the same beliefs about the need for and methods of solving joint problems, it is reasonable to expect that they will subsequently act on these beliefs. Cooperation, then, becomes the act of achieving a set of shared beliefs that are consistently labeled among the agents. Conflicts between agents show up as shared beliefs that are not consistently labeled, and negotiation is the process of relabeling beliefs in the appropriate agents. Exchanging information between rational agents can lead to agreement on the labels of shared beliefs if rational agents have enough information to understand each other's goals, intentions, and reasoning processes.

To implement this model for study, the Accord agent architecture employs justification–based reason maintenance systems as their belief system implementations. In this approach, each agent is regarded as possessing a graph of beliefs, B, that have labels giving their belief status, with the graph arcs giving justification information. Beliefs represent all the knowledge that an agent has, such as facts, agent positions, agent goals, agent intentions, and so forth. Beliefs can be private or shared with other agents, but all beliefs have a justification in at least one agent. Shared beliefs represent the information that agents must use to achieve cooperation. When working on problems, agents exchange information, and this shared information becomes shared beliefs. These beliefs contain all direct awareness that an agent has about the other agents. An extended version of the Huhns–Bridgeland algorithm (Huhns and Bridgeland, 1991) for distributed reason maintenance has been devised and a modular implementation developed for use by Accord agents in achieving a consistent multiagent belief system (Doulamis, 1994).

Multiagent communication is achieved through a message–passing technology. Each Accord agent is provided with a communication function to receive messages containing information and commands from outside the agent, and to send messages containing information and requests to outside agents. The bit–packet–level communication capability developed for this purpose is a system called Backtalk (Sayre & Gray, 1993). Backtalk consists of an agent server that exists to route messages among agents so that individual agents do not have to take this responsibility. It uses Unix pipes and sockets so that it can be used in any basic Unix environment. Backtalk does not provide knowledge representation formats for the messages it routes, and serves merely to route bit–strings, leaving the interpretation of the messages to the agents. It is similar to the communication capability implemented by Huhns, Bridgeland, and Arni (1990), but provides an agent server to allow dynamic addition and deletion of agents in support of open systems. The project is studying the use of KQML as a replacement for Backtalk.

An initial set of Accord agents have been implemented and used in building a simple multiagent system to perform routine von Neumann computer design. These agents use standard production–rule technology as their reasoning component, and the previously–mentioned multiagent reason maintenance system as their belief system. The agent set – a marketing agent, a processor designer, a memory designer, and a bus designer – produces a register–transfer–level design of a basic von Neumann computer. When conflict arises over contradictory design decisions, multiagent belief relabeling is used to arrive at a consistent design.

The Incorporation of Emotional Agents into The Accord Cooperation Mechanism

One of the limitations of the above line of work is the heavy dependence on the rationality of the agents. This approach can succeed if the agents are all computationally–based, but to develop practical, large–scale multiagent systems will necessarily involve bringing humans into these systems. This has caused the project to begin thinking how to include agents with emotional state in the cooperation mechanism.

This is especially important when considering how Accord agents can cooperate with human agents, since the rationality of humans is well–known to be strongly affected by their emotional state. The emotional state of a human agent strongly affects the conclusions drawn by human reasoning and the kind of cooperation taken as a result of this reasoning. Modeling emotional state will be fundamental to a computationally–based agent in its dealing with emotional agents.

A second justification for adding an emotional capability to a computationally–based agent is the possibility of using emotion as a particular sort of heuristic, filtering out extraneous information or ordering priorities, so that the agent is not burdened with too much information. Restricting the uncontrolled growth of belief systems is an important goal for justification–based reason maintenance research. While such a capability may be more remote at this point, it provides additional motivation for the work.

The project is currently studying the implementation of a model of emotional state based on Ortony, Clore, and Collins (1990) to investigate how emotional state might be taken into account during multiagent belief–based negotiation.

Elliott's (Elliott, 1994) model of emotional state, based on Ortony et. al., contains eight emotion groups, each containing two or more emotion types, for a total of 24 emotion types. Examples of these are (1) "joy : pleased about an event," (2) "resentment : displeased about an event desirable for another," (3) "hope : pleased about a prospective desirable event," (4) "disappointment : displeased about a disconfirmed desirable event," (5) "pride : approving of one's own act," and so forth. From this list, it is clear that cooperation between a computational–based rational agent and an emotionally–based human agent will require considerable understanding on the part of the computational agent. If a human agent is exhibiting joy after an event that is undesirable for a computational agent, the computational agent's approach to undoing the effects of this event must considerably different than if the human agent is exhibiting disappointment. In the disappointment situation, very little will be needed to convince the human agent that the situation needs to be changed, while in the joy situation, it may not be feasible to even attempt a change.

We analyze three approaches to including emotional state in multiagent reason maintenance. We note that even though it may be rational to have a certain belief state as a result of reason maintenance, the emotional state of an agent may cause it to act as if it had the opposite belief. The first way we are studying to account for this is by attaching lists of emotional states to beliefs, annotating the emotional states to inhibit or enable certain labels, forcing relabeling, thereby requiring the agent to exhibit a belief state consistent with its emotional state. A second way to approach this issue is by explicitly representing emotional state as beliefs, linking them through justifications to other beliefs, and organizing the justification network so that the belief statuses of the emotional state beliefs force a labeling consistent with the emotional state. The third way is by defining new gradations of agreement in multiagent reason maintenance. These new levels of consistency of the Huhns–Bridgeland type (Huhns and Bridgeland, 1991) would allow inconsistent labelings of beliefs to be acceptable within certain levels of agreement. Thus rational agents can understand that an agreement has been reached in which emotional agents may hold beliefs contrary to their rational interest.

The first acts more as a filter of label sets and thus does not slow down the relabeling process, but makes emotional state into an extralogical entity. The second slows down the relabeling process, but has the advantage that an agent's emotional state is explicitly represented in the belief base and can itself be reasoned about. The third has the advantage that emotional state is represented in a way that allows different levels of cooperation to be achieved without requiring complete agreement.

REFERENCES

Bridgeland, D. M. & Huhns, M. N. (1990). "Distributed Truth Maintenance." *Proceedings of AAAI–90: Eighth National Conference on Artificial Intelligence*, 72–77.

Doyle, J. (1979) "A Truth Maintenance System." *Artificial Intelligence*, 12, 231–272.

Doulamis, J. (1994) "A Distributed Truth Maintenance System." M. S. Thesis, The American University, Washington, DC.

Elliott, C. (1994). "Research problems in the use of a shallow Artificial Intelligence model of personality and emotion." *Proceedings of AAAI–94: Twelfth National Conference on Artificial Intelligence*, 9–15.

Genesereth, M. R. and Nilsson, N. J. (1987). *Logical Foundations of Artificial Intelligence,* Los Altos, CA: Morgan Kaufmann.

Huhns, M. N. and Bridgeland, D. M. (1991). "Multiagent Truth Maintenance," *IEEE Transactions On Systems, Man, and Cybernetics*, vol. 21, No. 6.

Huhns, M. N., Bridgeland, D. M., & Arni, N. V. (1990). "A DAI Communication Aide." *Proceedings of the 10th International Workshop on Distributed Artificial Intelligence*. Chapter 29.

Ortony, A., Clore, G. L., and Collins, A. (1990). *The Cognitive Structure of Emotions*, Cambridge, UK: Cambridge University Press.

Petrie, C., Russinoff, D., and Steiner, D. (1986). "Proteus: A Default Reasoning Perspective." *Proceedings of the 5th Generation Conference*, National Institute for Software.

Russinoff, D. (1985). “An Algorithm for Truth Maintenance.” Technical Report AI/TR–062–85, Microelectronics and Computer Technology Corporation, Austin, TX.

Sayre, K. and Gray, M. A. (1993). *“Backtalk: A Generalized Communication System for DAI.”* *Software - Practice and Experience*, Vol. 23(9), pp 1043-1057.